Philosophy Compass 1/2 (2006): 201–223, 10.1111/j.1747-9991.2006.00020.x

Fast and Frugal Heuristics[1]

Michael A. Bishop
Northern Illinois University

Abstract
A heuristic is a rule of thumb. In psychology, heuristics are relatively simple rules for making judgments. A fast heuristic is easy to use and allows one to make judgments quickly. A frugal heuristic relies on a small fraction of the available evidence in making judgments. Typically, fast and frugal heuristics (FFHs) have, or are claimed to have, a further property: They are very reliable, yielding judgments that are about as accurate in the long run as ideal non-fast, non-frugal rules. This paper introduces some well-known examples of FFHs, raises some objections to the FFH program, and looks at the implications of those parts of the FFH program about which we can have some reasonable degree of confidence.

Part I: Introduction

A heuristic is a rule of thumb. In psychology, heuristics are relatively simple rules for making judgments. A fast heuristic is easy to use and allows one to make judgments quickly. A frugal heuristic relies on a small fraction of the available evidence in making judgments. Typically, fast and frugal heuristics (FFHs) have, or are claimed to have, a further property: They are very reliable, yielding judgments that are about as accurate in the long run as ideal non-fast, non-frugal rules.[2] FFHs are most fully explained in *Simple Heuristics that Make Us Smart* (Gigerenzer, Todd, and the ABC Group 1999).

This paper is organized as follows. Part I introduces some background and two well-known examples of FFHs. Part II raises three objections to the FFH program. While I think these challenges are serious, given the relative youth of the FFH program, it's probably too early to say whether they can be overcome. Part III articulates some potential implications of those parts of the FFH program about which we can have some reasonable degree of confidence.

Before proceeding, I should issue a warning: My primary goal here is not to settle any issues – although I will certainly try! My primary goal is to clearly articulate the FFH program, some of its challenges, and some of its possible implications. I want readers to come away convinced that the FFH program raises issues that are philosophically interesting and important – even if (especially if!) they ultimately reject my arguments.

1. BACKGROUND

It is useful to think about the FFH program arising in response to two important literatures in psychology which, when stripped of lots of detail and complexity, offer the following lessons:

1. Heuristics and Biases: (a) We are naturally disposed to use simple reasoning rules, and (b) these rules are less reliable in the long run than ideal rules (e.g., Kahneman, Slovic, and Tversky 1982, see also Simon 1982).
2. Predictive Modeling: (a′) There are simple reasoning rules we don't use, and (b′) these rules are about as reliable in the long run as ideal rules (e.g., Grove and Meehl 1996; Breiman, *et al.* 1984).

These lessons are prima facie pessimistic about our reasoning abilities. The FFH literature embraces half of each of these pessimistic lessons, for an optimistic result:

3. Fast and Frugal Heuristics: (a) We are naturally disposed to use simple reasoning rules, and (b') these rules are about as reliable in the long run as ideal rules.

Let's look in some detail at arguably the two most important FFHs, Take the Best and the recognition heuristic.

2. TAKE THE BEST

Suppose you must predict which of two cities has the higher rate of homelessness on the basis of six cues (see Chart 1). The trick, however, is that your evidence is imperfect – all you know is in what direction this evidence "points" (Gigerenzer, Czerlinski, and Martignon 2002, p. 562). One cue is naturally binary (rent control) and the others are coded in terms of whether the city is above or below the median (for the largest 50 cities in the U.S.) on the particular cue. A "1" is assigned to a city when the cue suggests a higher rate of homelessness, otherwise a "0" is assigned. When a cue assigns a value of "1" to one city and a value of "0" to another, the cue *discriminates* between the cities. So, in Chart 1, the *rent control* cue discriminates between Los Angeles (1) and Chicago (0). The cues are ordered by their validities – a cue's validity is its percentage of right answers in those cases in which it discriminates between cities.[3] Given the information in chart 1, how would you predict which of two cities has the higher rate of homelessness?

Take the Best (TB) begins with the cue with the highest validity. If the cue discriminates, then TB predicts that the object assigned "1" has the higher value; if it doesn't discriminate, then TB continues on to the cue with the next highest validity. And so on. If no cue discriminates, TB guesses (in practice, TB is assumed to get 50% of these cases correct). TB is frugal. There are six possible city pairings in Chart 1. In four, TB considers only one cue; in one (Chicago-New Orleans) it considers two cues; and in only one

Chart 1. The Homelessness Problem

	Los Angeles	**Chicago**	**New York**	**New Orleans**
Rent control (1 is yes)	1	0	1	0
Vacancy rate (1 is below median)	1	1	1	0
Temperature (1 is above median)	1	0	1	1
Unemployment (1 is above median)	1	1	1	1
Poverty (1 is above median)	1	1	1	1
Public housing (1 is below median)	1	1	0	0

Adapted from Gigerenzer, Czerlinski and Martignon (2002, p. 562)

does it consider all the cues (Los Angeles-New York). Compare TB to a *unit model*, which adds up the columns and predicts the object with the highest sum (and guesses if they're equal). The unit model always considers all six cues in making its prediction, and so TB is more frugal than the unit model. TB is also faster than the unit model – and not just because it is more frugal. The unit model integrates the different lines of evidence by adding together the scores. TB doesn't. As it happens, the cities are arranged in terms of their homelessness rates, from highest to lowest. So of the six possible pairings in Table 1, TB and the unit model make the exact same judgments and get five pairings right. (They're wrong about New York-Chicago.)

How accurate is TB? This is a harder question than it seems. Consider that in order to construct a TB model, we need to know the validities of each cue (that is, the percentage of right answers each cue gives in those cases in which it discriminates between cities). The TB model needs some data to "learn" what the cue validities are. Call this data the *training set*. When the TB model is tested, it will be given some data and then asked to make a prediction (e.g., Which city has the higher rate of homelessness?). Call the data on which the model is tested the *test set*. There are two ways psychologists typically test models against each other:

(a) Curve-fitting test: The training set is identical to the test set. So the models "learn" from the same data on which they are tested.
(b) Prediction test (aka cross-validation): The training set is different from the test set. So the models "learn" from one set of data and are tested on a different set of data.

Prediction tests can be done in a number of different ways. Gigerenzer, Czerlinski, and Martignon randomly divide the 50 largest cities into two sets of 25 cities. The 300 possible pairings in the first set of 25 cities are used as the training set; and the 300 possible pairings in the second set are used as the test set. So a model that gets 210 right out of the 300 test set pairings gets a reliability score of 70% on those pairings. This procedure (with different random divisions into training and test sets) is repeated 1,000 times to determine a model's reliability score (2002, p. 568).

Gigerenzer, Czerlinski, and Martignon tested TB against the unit model. In the training set, the unit model must "learn" the direction of the cues. For example, if unemployment is positively correlated with homelessness in the training set, then any city in the test set that is above the median in homelessness gets a "1" and any city that is below the median in homelessness gets a "0." TB was more reliable than the unit model on both curve-fitting and prediction tests (2002, p. 568). But how would TB fare against stiffer competition? Gigerenzer, Czerlinski, and Martignon (2002) tested TB against a multiple regression model and a Bayesian network. Here is a brief and intuitive introduction to these types of models.

1. Multiple Regression. An intuitive way to understand multiple regression models is to suppose we are considering a case in which we have a single cue (a person's height) and a target property we want to predict (the person's weight). Suppose we take a training set (the heights and weights of a large number of people) and plot them on a graph. A regression model would draw a straight line on the graph that comes closest to all the data points. In most realistic cases, however, we will have more than a single predictor cue, and so we need a multiple regression equation. Such equations have the form:

$$y = k + c_1x_1 + c_2x_2 + c_3x_3 \ldots$$

The trick to building a regression model involves choosing the coefficients (c_n) (aka the weights) so that the model best fits the data in the training set.

2. Bayesian networks. At the heart of any Bayesian network is a commitment to update the probability of a state or event in terms of Bayes theorem. To see how Bayes' theorem works, suppose that there is a 10% chance that S has a disease, D. This is the prior probability of D (or P(D)). Now suppose that we get a new piece of evidence, E: S tests positive for D. We want to know what the probability is that S has D given her new evidence, E. This is the posterior probability of D (or P(D / E)). In this situation, Bayes' theorem tells us to update our probabilities according to the following equation:

$$P(D / E) = \frac{P(D) \times P(E / D)}{[P(D) \times P(E / D)] + [P(-D) \times P(E / -D)]}$$

In order to figure out the posterior probability (the probability that S has D given the positive test result), we must know the test's sensitivity and specificity. The sensitivity of the test (aka the likelihood of the evidence) is the probability that a person with the disease will get a positive result on the test: P(E / D). The specificity of the test for D is the probability that a person without the disease will get a negative result: P(–E / –D). (Readers will note that Bayes theorem contains two probabilities I haven't mentioned: P(–D) and P(E / –D). According to the laws of probability, the first can be derived from the prior probability, P(–D) = 1 – P(D), and the second can be derived from the test's specificity, P(E / –D) = 1 – P(–E / –D).) So suppose that the test's specificity and sensitivity is 90%. Plugging in the various values, we find that the probability that S has D given the positive test result is ½.

$$P(D\ /\ E) = \frac{.1 \times .9}{(.1 \times .9) + (.9 \times .1)}$$

A simple Bayesian network consists of nodes, which represent states of affairs, and their probability relations. For example, Charniak describes a Bayesian network that assigns a probability to whether your family is home on the basis of whether the outdoor light is on (which is usually left on when everyone is out) and whether the dog is out (which usually happens when everyone is out or the dog has a bowel problem); and you can determine whether the dog is out on the basis of whether you hear barking (1991, pp. 50–1).

The probability distribution for this network is defined by the prior probabilities of the root nodes (those with no predecessors – the top two nodes) and by the conditional probabilities of each of the other nodes given all possible combinations of their immediate predecessors. As new evidence comes in, the conditional probabilities of the nodes can be updated according to Bayes' theorem. So this network can calculate the conditional probability that your family is out given that the light is on and you don't hear barking. Bayesian updating on the above network is relatively easy (for a computer!) since there aren't a lot of nodes and (more importantly) there is just one path connecting any two nodes. But for more complex problems, and particularly when nodes are multiply connected, Bayesian networks become computationally intractable. In recent years, there has been a proliferation of networks that can handle complex problems

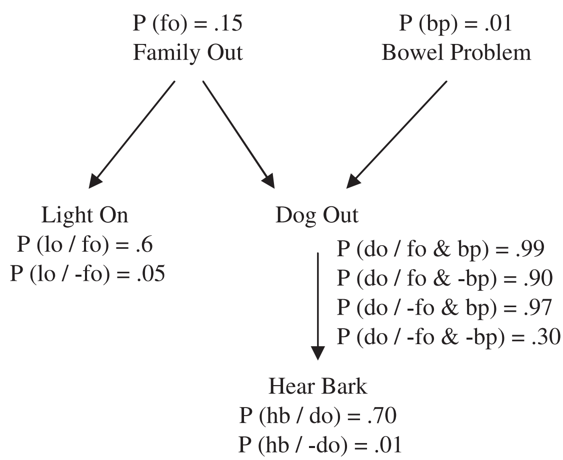

Figure 1. A Bayesian network (adapted from Charniak 1991, 52).

by finding approximate Bayesian solutions. In fact, Gigerenzer, Czerlinski, and Martignon employ just such a network (2002, pp. 575–7).

Although we have not fully delved into the details of the regression and Bayesian models, it should be clear that they are much more complex than TB. So how do they compare in terms of their accuracy? Gigerenzer, Czerlinski, and Martignon (2002) compared TB to multiple regression and Bayesian models on four real-world data sets.

Let's start with the first curve fitting tests. The most reliable curve-fitting model is the Profile Memorization Method (PMM). It memorizes the ordering of the profiles for all 50 cities, but it doesn't memorize the cities' names. So it remembers that a city with the profile <1, 1, 1, 1, 1, 1> has a higher homelessness rate than a city with the profile <0, 1, 0, 1, 1, 1>. But it doesn't remember that Los Angeles (a city with the first profile) has a higher homelessness rate than Chicago (a city with the second profile); if it remembered the cities' rankings, it would never make a mistake. The PMM makes mistakes because different cities can have the same profile. Given a prediction task involving a profile shared by more than one city, the PMM makes the prediction that is most likely to be correct or, if neither is more likely to be correct, it guesses (Gigerenzer, Czerlinski, and Martignon 2002, p. 568). On the four real-world data sets, the Bayesian network outperformed TB by between 2% and 8% on the curve-fitting tests (chart 2). On the prediction tests, the differences closed dramatically. TB was within 1% to 2% of the Bayesian network (chart 3).[4]

Martignon and Laskey (1999) ran the same sort of tests on 20 wide ranging real-world data sets (see chart 4).

Chart 2. Results of Curve-Fitting Tests

	Take the Best	Mult Regression	Bayesian Network	Profile Mem Meth
City population	74	74	76	80
Homelessness	69	70	77	82
Fish fertility	73	75	75	75
Profs' salaries	80	83	84	87

Adapted from Gigerenzer, Czerlinski and Martignon (2002, p. 578)

Chart 3. Results of Prediction Tests

	Take the Best	Mult Regression	Bayesian Network
City population	72	71	74
Homelessness	63	61	65
Fish fertility	73	75	75
Profs' salaries	80	80	81

Adapted from Gigerenzer, Czerlinski and Martignon (2002, p. 579)

Chart 4. Results of Tests on 20 Real-World Data Sets

	Take the Best	Mult Regression	Bayesian Network	Profile Mem Meth
Curve-fitting	75	77	79	85
Prediction	71	68	74	

Adapted from Martignon and Laskey (1999, p. 182)

In my view, Take the Best is the most impressive heuristic to come out of the FFH literature. Before turning to the issue of just how impressive it is, though, let's consider a second FFH.

3. THE RECOGNITION HEURISTIC

You have to be a bit ignorant to use the recognition heuristic. Consider the question: Which city has more inhabitants, San Diego or San Antonio? Goldstein and Gigerenzer found that while U.S. students got the answer right 62% of the time, German students got it right 100% of the time (1999, p. 43). To see if this was mere luck, they randomly paired the 22 largest U.S. cities, and they randomly paired the 22 largest German cities, and asked U.S. students to pick the most populous. The U.S. students were more reliable about the German cities than the American cities (pp. 53–4). Goldstein and Gigerenzer call this the *less-is-more effect*: Under certain circumstances, greater ignorance can yield greater reliability. The idea is that subjects who are somewhat more ignorant about a topic are more likely to employ the *recognition heuristic*: If S recognizes one of two objects but not the other, and recognition correlates positively (negatively) with the criterion, then S can infer that the recognized object has the higher (lower) value.[5] And when the subject is right about the correlation, the recognition heuristic can be more reliable than slower, more expensive reasoning strategies.

In a striking paper, Borges, Goldstein, Ortmann, and Gigerenzer (1999) apply the recognition heuristic to the area of investment. They test a number of different investment strategies against the recognition heuristic (investments were selected on the basis of name-recognition). For the six-month period of study, the recognition heuristic outperformed the other strategies (Borges *et al.* 1999, p. 65). Borges *et al.* suggest that that ordinary people (who are disposed to naturally use the recognition heuristic) can perhaps do better on the stock market than mutual fund managers and market indices: "In investments, there may be wisdom in ignorance" (1999, p. 72).

Part II: Some Challenges

The research on FFHs is fascinating and prima facie impressive. But it is a relatively new program, and not without challenges. I would be remiss not to mention at least some of these challenges. But I won't pretend to offer definitive judgments about whether they can be overcome. Much will depend on the outcome of future research. My goal here will be to introduce loose ends. Tying them up will be left for other occasions.

1. ARE PEOPLE NATURALLY DISPOSED TO USE FFHS?

The FFH literature would lose much of its distinctive appeal if our use of highly reliable FFHs were very occasional – that is, if we almost never reason in ways that are fast, frugal, and close to optimally reliable. That's because we already know from the predictive modeling literature that there are simple, highly reliable reasoning rules. What's distinctive about FFHs is the idea that we naturally employ such rules as our default reasoning strategies. The evidence for thinking that we are naturally disposed to use FFHs is real (Rieskamp and Hoffrage 1999), but spotty. The less-is-more effect suggests that people are naturally disposed to use something like the recognition heuristic (together, perhaps with the availability heuristic – see endnote 4). Further, Payne, Bettman, and Johnson (1988, 1993) present evidence that people favor strategies like Take the Best when they are under time constraints. More research is needed to confirm that we have a consistent tendency to employ FFHs.

2. HOW ROBUST ARE FFHS?

FFHs are "domain specific." The recognition heuristic, for example, "only works in environments where recognition is correlated with the criterion" (Goldstein and Gigerenzer 1999, p. 41).[6] But there is a tradeoff between domain specificity and robustness – a highly domain specific rule will, by definition, not perform well across a wide range of circumstances. A lack of robustness need not be a problem if people can easily recognize and apply the heuristic to its appropriate domain. But there is reason to doubt that people are especially good at this. For example, the recognition heuristic did very well on the stock market over a six-month period (Borges *et al.* 1999). Is this because the heuristic is robust (and so can be expected to perform well in a wide range of market environments) or lucky (just happening to be successful in that six-month period)? We don't know – and so we don't know whether it's reasonable to apply it to the stock market.[7] Further, Take the Best seems to be reliable in rather non-obvious and complex conditions (Martignon and Hoffrage 1999). So we might admit that FFHs are in principle highly reliable. But given the difficulties people are likely to have properly applying them, in the real world FFHs might generally not be anywhere near as reliable as advertised. (Many commentators on Todd and Gigerenzer's target BBS article (2000) discuss this problem.)

3. HOW RELIABLE ARE FFHS, REALLY?

Some of the results in Charts 2–4 are breathtaking. But keep in mind that the various rules are being tested on binary problems. So there is a reliability floor of 50%. What's more, the reliability ceiling is less than 100%. In the curve fitting tasks, the Profile Memorization Method did no better than 87%, and in the prediction tasks, the Bayesian network did no better than 81%. So consider that TB is 3% less reliable in prediction tasks than the Bayesian network in Chart 4. Sounds impressive. But if we assume that the Bayesian network is at, or close to, the reliability ceiling, this means that TB is about 10% worse than the Bayesian network *within the range of possible reliability scores*. Perhaps not quite so impressive. (Payne, Bettman, and Johnson (1993, pp. 88–91, 128–33) argue that this is the right way to understand the relative reliability scores of different predictive models.)

Part III: Some Potential Implications

Having reviewed some of the challenges facing the FFH program, one might wonder whether there is anything positive – much less anything of philosophical interest – to extract from it. We should recognize up front that doing philosophy in a way that takes empirical findings seriously is risky. But philosophy is risky. It is not clear that it is less risky to do philosophy uninformed by science. And anyway, what's so great about non-risky research? Risky research can bring greater potential for rewards than safe research. Finally, the very real risk of relying on bad empirical theories can be somewhat allayed if we can identify elements of a promising theory about which it is reasonable to have quite a lot of confidence (Kitcher 1993). Looking at the evidence, I think there are at least two elements of the FFH program about which we can have a good deal of confidence.

1. People use heuristics in their reasoning.
2. There are simple heuristics that if appropriately applied are (close to) optimally reliable.

Evidence for these two propositions comes from the FFH program. But as I mentioned in the introduction, evidence for (1) also comes from the heuristics and biases program, and evidence for (2) also comes from the predictive modeling literature. There is a third element of the FFH program for which there is at least some evidence (see Part II, section 1):

3. On occasion, people use heuristics that are (close to) optimally reliable.

As I have noted, the evidence for (3) is spotty. Still, we can extract positive lessons from (1) and (2) if it should turn out that (4) (which is not incompatible with [3]) is true:

4. People can often learn to use heuristics that are optimally reliable (or close to it).

There is some evidence that a relatively modest amount of training can help people to reason better (e.g., Richard Nisbett 1993).

In my view, there are enough well-documented findings here for it to be reasonable to think about the philosophical assumptions and implications of the FFH program. Close attention to the FFH literature raises a host of issues in the philosophy of psychology. For example, I have described heuristics as relatively simple rules for making judgments. But Gigerenzer argues that heuristics possess three further features: They must (a) exploit evolved capacities, (b) exploit structures of environments, and (c) describe actual psychological processes (2004, pp. 63–4). Each of these conditions raises issues about the nature of mind and how to properly study it. (a) raises questions that have received a lot of attention recently about the role of evolution in shaping our minds (Barkow, Cosmides and Tooby 1992, Buller 2005). (b) relates to the "domain specificity" of our heuristics (see Endnote 5). It also raises the question of whether our minds are "massively modular" (Barkow, Cosmides and Tooby 1992, Pinker 1997, Samuels 1998, Fodor 2000). And implicit in (c) is a methodological constraint on accounts of heuristics. Gigerenzer demands computational models of heuristics: A properly described heuristic "can be instantiated as a computer program" (Gigerenzer and Todd 1999, p. 16). These are important issues, and spelling them out could easily take up the rest of this paper. But having pointed them out, I now propose to abandon them. Instead, I want to briefly sketch what I take to be some provocative – but far from established – epistemological implications of the parts of the FFH program about which it is reasonable to be confident.[8] The following four sections will address the following issues:

1) The three literatures described here – the heuristics and biases program, the FFH program and the predictive modeling literature – all make explicit normative, epistemological prescriptions about how we ought to reason and what we ought to believe. In a nutshell, they are explicitly normative areas of psychology.
2) The attempt to articulate the normative framework that grounds the prescriptions made by the FFH literature is an interesting and legitimate epistemological endeavor. In rough terms, this framework involves the assessment of reasoning strategies in terms of their reliability and tractability.
3) The normative nature of this psychological research makes it inevitable that psychologists sometimes get embroiled in normative debates. I suggest a partial resolution of one such debate.
4) There is some reason to prefer the epistemological framework of the FFH program to the theories of justification defended by contemporary epistemologists.

The following discussion is animated by my belief that the psychological literatures described here are rich sources of epistemological theorizing that have been too much ignored. But these are not meant to be the last words on any of these subjects. They are put forward as suggestive and preliminary first words.

1. PSYCHOLOGY AND EPISTEMOLOGY

All branches of science tell us how we ought to reason. Physics tells us how we ought to reason about physical matters, chemistry tells us how we ought to reason about chemical matters, etc. But the three areas of psychology we have considered (predictive modeling, heuristics and biases, and fast and frugal heuristics) are different. They make explicit judgments about how we ought to reason about matters that have nothing to do with the subject matter of psychology. If a European wants to know whether San Diego is more populous than San Antonio, she ought to use the recognition heuristic. This seems to be an *epistemological* judgment about how someone ought to reason about demography. This kind of normative, epistemological judgment is common in these literatures. The title of Gigerenzer and Hoffrage's 1995 article explicitly describes their normative aspirations: "How to *Improve* Bayesian Reasoning Without Instruction" (emphasis added). The heuristics and biases literature is famous (and in some circles, infamous) for passing judgments about how badly people reason (e.g., Piattelli-Palmarini 1994). There is considerable debate about whether we should accept these pessimistic judgments about our reasoning (e.g., Cohen 1981, Gigerenzer 1991). Turning to the third body of research, the predictive modeling literature, we find literally scores of studies that deliver advice about how we *ought* to reason about a wide variety of topics. These topics include judgments about violence, criminal recidivism, preliminary psychiatric diagnoses, academic performance, credit risk, susceptibility to Sudden Infant Death Syndrome, the quality of the vintage for a red Bordeaux wine, the presence, location and cause of brain damage, whether someone is suffering from progressive brain dysfunction, and whether electroshock therapy will be successful for a psychiatric patient (for citations, see Bishop and Trout 2005, pp. 13–14).

These areas of psychology are chock full of normative, epistemological judgments. What's more, powerful institutions have followed some of this advice (e.g., financial institutions and hospitals). In order to make such judgments, these areas of psychology must be presupposing some sort of general epistemological framework. A project naturally suggests itself: Articulate the epistemological framework that grounds the normative judgments of these areas of psychology (Bishop and Trout 2005). For those who might have worries about whether this is a reasonable project, we should ask: Reasonable compared to what? I would argue that this project makes at least as much sense as the standard project in contemporary analytic epistemology, which seems to involve articulating an epistemological framework that captures our intuitive judgments about knowledge and justification (Weinberg, Nichols, and Stich 2001). The standard epistemological framework aims to capture our intuitive judgments, whereas the proposed framework aims to capture the judgments of these successful areas of psychology. To make the case for the legitimacy of the latter framework, let's note that in many important situations, the judgments delivered by psychological research are better than the judgments delivered by our unaided intuition. In what sense better? In the long run, following that advice will result in our adopting more true and useful beliefs. For example, the Violence Risk Appraisal Guide (VRAG) is more reliable than experienced parole officers in making judgments about whether a felon will commit a violent crime if paroled (Quinsey *et al.* 1998). Why not prefer the normative framework that captures the better judgments about how we ought to reason and what we ought to believe about some very high-stakes issues? It is hard to see what sort of *objection* one could have to this project (besides, perhaps, that it will inevitably result in an uninteresting, pedestrian epistemological theory – see Part III, Section 4). If anything, this (admittedly contentious) way of putting the issue might raise worries about how reasonable it is to pursue the traditional project. I won't press that point here.[9] My goal is simply to motivate this new project.

Before moving on, let's briefly reflect on the nature of this proposed epistemological project. In "Epistemology Naturalized" (1969), Quine famously argues that epistemology is a branch of psychology. The received wisdom holds that Quinean naturalism is subject to a devastating objection: It empties epistemology of its normative character. This objection rests on a common view about science and philosophy. Science is entirely descriptive, while philosophy is the exclusive and proper domain of the normative. This picture leads many philosophers to insist that science is irrelevant to normative theorizing – or perhaps relevant only at the margins if we assume that ought implies can. But is psychology really empty of normative content? I have suggested that three areas of psychology make lots of normative judgments, and I have further suggested that there is some interesting epistemology to be done by exploring the framework that grounds these normative judgments. If this is right, if parts of psychology are normative, then epistemology can be a chapter of psychology while maintaining its status as a normative discipline. So maybe Quine was right.[10]

2. THE NORMATIVE FRAMEWORK OF THE FFH PROGRAM

Any epistemological theory will aim to evaluate some cognitive items in terms of some normative category. Most contemporary epistemological theories aim to set out the conditions under which a belief token is justified or is an instance of knowledge: Is S's belief that p justified? Does S know that p? The FFH literature has a very different orientation. The primary items of evaluation are not belief tokens but reasoning strategies. In the research we reviewed in Part I, the items being evaluated included Take the Best, Unit Weight Models, Multiple Regression Models, Bayesian

Models, the Profile Memorization Method, and the Recognition Heuristic. These are not belief tokens. They are methods or strategies for reasoning about a certain class of problem. It is the same in the predictive modeling literature. The primary objects of study, and the objects that get evaluated in the first instance, are reasoning strategies.[11] How are reasoning strategies evaluated? In the FFH and predictive modeling literatures, they are evaluated in terms of their *reliability* (i.e., their tendency to produce truths) and their *tractability* (i.e., how easy they are to use). So other things being equal, a more reliable reasoning strategy is superior to a less reliable one, and a fast and frugal heuristic is superior to a slow and informationally demanding reasoning strategy. The FFH and the predictive modeling literatures both presuppose a reliabilist framework for evaluating reasoning strategies (Bishop 2000, Bishop and Trout 2005). But in the case of the heuristics and biases program, the issue of how to evaluate different reasoning strategies – different heuristics – has a more complicated history. And it is a history that has become entwined with that of the FFH program.

3. Normative Debates in Psychology

A major characteristic of the heuristics and biases program (HB) is its focus on people's systematic reasoning errors (or biases). This focus has led to considerable debate about whether these biases really are "errors" and whether, or to what extent, people are rational (e.g., Nisbett and Borgida 1975, Cohen 1981, Stich 1985, 1990, Lopes 1991, Gigerenzer 1991, 1996, Kornblith 1992, Piattelli-Palmarini 1994, Stein 1996, Sosa and Galloway 1999). This is a normative debate, and a central issue is: What makes something a bias or a systematic reasoning error? According to Kahneman and Tversky, "[t]he presence of an error of judgment is demonstrated by comparing people's responses either with an established fact . . . or with an accepted rule of arithmetic, logic, or statistics" (1982, p. 493). This is not a reliabilist answer of the sort we find in the FFH or the predictive modeling literature. It harkens back to the traditional program in epistemology in that it evaluates particular *belief tokens*. To see how this view works, consider the following diagnosis problem.

> The probability of breast cancer is 1% for women at age forty who participate in routine screening. If a woman has breast cancer, the probability is 80% that she will get a positive mammography. If a woman does not have breast cancer, the probability is 9.6% that she will also get a positive mammography. A woman in this age group had a positive mammography in a routine screening. What is the probability that she actually has breast cancer? __%.

The Bayesian answer is about 7.8%. Gigerenzer and Hoffrage report that 16% of subjects faced with this problem got the Bayesian answer (1995, p. 693). So for proponents of HB, 84% of subjects made an error because their response violated an accepted rule of statistics, namely, Bayes' theorem.

Have these subjects made an error? And if so, what is the nature of their error? It is important to keep these two issues separate. We might agree with proponents of the HB program that these subjects have made an error, but we might disagree with them about the nature of this error. A rather heated debate about these normative issues broke out between proponents of the HB and the FFH programs (see Kahneman and Tversky 1996, Gigerenzer 1991, 1996).

Gigerenzer, the main proponent of FFHs, levels two objections against the HB program (1991). Both objections focus on examples like the Bayesian problem above, in which subjects are asked to judge the likelihood of an individual event.

1) *The Disappearance Argument*. The errors subjects make on the diagnosis problem can be made to "disappear" (or at least be substantially reduced) by framing the problem differently. Gigerenzer and Hoffrage (1995) gave a group of subjects a problem that is mathematically equivalent to the Bayesian problem, except it was framed in terms of frequencies rather than probabilities. On this version of the problem, 46% of subjects with no training in statistics got the Bayesian answer (693).
2) *The Frequentist Argument*. A frequency interpretation of probability states that the probability of an attribute A is the relative frequency with which A occurs in an unlimited sequence of events. According to a frequentist, assigning a probability to a single event is meaningless. So any answer to the meaningless question posed by the above problem ("What is the probability that she has breast cancer?") cannot be a violation of probability theory. As we have already seen, Kahneman and Tversky contend that "[t]he presence of an error of judgment is demonstrated by comparing people's responses either with an established fact . . . or with an accepted rule of arithmetic, logic, or statistics" (1982, p. 493). Since a frequentist would not take subjects' answers to violate any "rule of arithmetic, logic, or statistics," those answers are not errors.

These are very puzzling arguments. The second argument aims to show that people aren't making (and in fact, can't make) errors on single event probability problems. But the first argument shows how people can *improve* their reasoning on those exact problems (i.e., by reframing them in terms of frequencies). The second argument contends there is no error, while the first argument shows how to correct the error. Samuels, Stich, and Bishop argue that Gigerenzer can't have it both ways: "If it ain't broken, you *can't* fix it" (2000, p. 250).

To make matters worse, the Frequentist Argument involves at least three anomalies. First, Gigerenzer never suggests that subjects understand these questions in a frequentist way. But he criticizes the HB program for not paying attention to precisely this issue, namely, how subjects understand these problems (1996, p. 593). Second, Gigerenzer insists that the point of the Frequentist Argument "is precisely *not* to champion" a frequentist interpretation of probability (1994, p. 141). But if he doesn't accept a frequentist interpretation of probability, why does the Frequentist Argument insist on it? And third, essential to the Frequentist Argument is the HB view about what is involved in a reasoning error – violating a law of probability. But Gigerenzer is unequivocal in his rejection of this view. He clearly embraces a broadly reliabilist approach to epistemic evaluation:

> We do not compare human judgment with the laws of logic or probability, but rather examine how it fares in real-world environments. The function of heuristics is not to be coherent. Rather, their function is to make reasonable, adaptive inferences about the real social and physical world given limited time and knowledge. (Gigerenzer and Todd 1999, p. 22)

Given these anomalies, we might charge Gigerenzer with confusion or worse (as critics have done, see Kahneman and Tversky 1996, Samuels *et al.* 2002, Bishop and Trout 2005). But these anomalies point to a different reading of the Frequentist Argument. I propose an interpretation that smoothly handles the anomalies, but that Gigerenzer (to my knowledge) has never clearly and explicitly embraced.

The only way Gigerenzer's Frequentist Argument makes sense is if we understand it to be a *reductio ad absurdum*. Here are the words I propose to put into Gigerenzer's mouth:

> Suppose we adopt the HB standard of what counts as an error – a violation of the rules of logic or probability. Unless proponents of the HB program have an argument against the frequentist interpretation of probability (and they don't), they have no way to block the committed frequentist's argument. And so they can't legitimately conclude that the subjects are making reasoning errors. But this is absurd – those subjects *are* making reasoning errors. In fact, there are some nifty ways of getting people to make fewer such errors (see the Disappearance Argument). Therefore, in order to avoid this absurd conclusion, we must reject either frequentism or the HB standard of what counts as an error. Which disjunct will the proponent of the HB program reject? In perhaps the most stinging criticism of my [Gigerenzer's] research, Kahneman and Tversky insist that the HB program has relied heavily on studies (on availability, anchoring, and overconfidence) in which probability statements were interpreted as frequency claims (1996, pp. 583–4). So don't take it from me. Take it from Kahneman and Tversky: The only disjunct the proponent of the HB program can reject is the second – the HB standard of what counts as an error.

Charity demands that we interpret the Frequentist Argument as a *reductio*. While this interpretation saves Gigerenzer from charges of confusion (or worse), it does not follow that the Frequentist Argument is sound. (For an interesting exchange on this issue, see Vranas 2000, 2001 and Gigerenzer 2001).

4. THE STATUS OF THE NORMATIVE FRAMEWORK OF THE FFH PROGRAM

The epistemological framework that grounds the normative judgments of the FFH program evaluates reasoning strategies in terms of their reliability and tractability. This could be sharpened considerably, but for the purposes of our discussion here, let's rely on this rough characterization. Is there any reason to take this framework seriously – seriously enough (say) to develop it in more detail? Perhaps. But a natural worry is that the reasonable prescriptions of the FFH program (and of the HB program and the predictive modeling literature) can be co-opted by any plausible standard epistemological theory. This "co-opting strategy" is bound to be appealing to many epistemologists. It doesn't involve denying any empirical findings or rejecting any of the reasonable normative evaluations made by psychologists. Rather, it involves showing that when we grant psychologists their empirical claims, there's nothing particularly surprising or interesting about their normative claims. For example, the co-opting epistemologist will grant (at least for the sake of argument) that TB is highly reliable and that this descriptive result is fascinating and important. But once the descriptive result has been granted, the co-opting epistemologist will argue that the resulting normative claims are utterly pedestrian: If S is naturally disposed to use TB, the beliefs that result from her doing so are epistemically justified. Ho hum.

I think the "co-opting strategy" is problematic. The FFH prescriptions are not pedestrian; in fact, some are so counterintuitive, I suspect that some epistemologists won't *want* to co-opt them. In what follows, I will argue that (1) a theory of justification that co-opts the FFH program must accept some judgments that are not obviously intuitive, (2) any theory that absorbs the FFH program and accepts these not obviously intuitive judgments will face a serious objection, the Generality Problem, and (3) the epistemological framework that grounds the prescriptions of the FFH program neatly avoids the Generality Problem.

4.1. *The Challenge of Co-Opting the Results of the FFH Program*

The epistemological framework that grounds the FFH program recommends belief-forming strategies. So suppose S learns that an easy-to-use heuristic (say, TB) is very reliable – in fact, more reliable than she is on a certain class of problem. (This assumes that S is not naturally disposed to use TB on this class of problems.) She learns to apply TB, and so she applies it to a particular problem and comes to believe that p. S knows that by using TB in this case, she is ignoring a lot of readily available evidence. S does not, however, know that had she considered all the available evidence, she would *not* have accepted the belief that p, the belief recommended by TB. She would have recognized that the evidence actually supports a different belief, the belief that r. Let's suppose, further, that r is true. In this case, let's say that the FFH led S to an "I-error" – a false belief that S could have avoided had she considered all the evidence.[12] To see that I-errors are bound to happen, suppose TB recommends the belief that p on the basis of a single cue, but the rest of the cues indicate that r is true. Being savvy about such matters, S knows that I-errors are inevitable if she repeatedly uses TB. But given any particular application of TB, she doesn't know whether it is leading her to make an I-error.

What justifies S in accepting a belief that results from her using TB, particularly since she knows that TB will sometimes lead her to make I-errors? She can argue as follows:

> I know that TB is more reliable than my unaided reasoning. So I will sometimes be disposed to disagree with TB. But when I disagree with TB, TB will be more often right. (Otherwise, it would not be more reliable than I am.) Of course, sometimes I will be right and TB will be wrong, and in those cases following TB will lead me to make an I-error. But in the long run, while my use of TB will lead to more I-errors, it will lead to fewer overall errors.

Let's focus on a particular I-error. S has used TB and as a result has accepted a false belief; had S considered all the evidence available to her, she would have reasoned to the true belief that was best supported by that evidence. Is S's belief (the I-error) justified? One might reasonably judge that it is *not* justified. But then one's theory of justification cannot co-opt the normative recommendations of the FFH program. *The entire point of the FFH program is that people can, and often should, use very reliable FFHs that ignore lots of evidence and do not properly integrate the evidence they do consider*. Following this advice will inevitably lead to some I-errors. The epistemological framework that grounds the FFH program sanctions these I-errors for the sake of fewer overall errors. A standard theory of justification can co-opt these recommendations only if it is willing to admit that these occasional I-errors are epistemically justified.

My suspicion is that at least some proponents of standard theories of justification will insist that in the situation described above, the I-error is not epistemically justified. If so, then these epistemologists cannot (and do not want to) co-opt the FFH results. But consider what is lost on this view. We have already noted that there are lots of simple rules for making important judgments – including medical diagnoses and predictions about violence. Does the proponent of a theory of justification really want to say that a parole officer is not justified in accepting beliefs on the basis of the VRAG, because the VRAG will occasionally lead to I-errors (even though it will lead to fewer overall errors)? Seems like a high price to pay.[13]

4.2. *Co-Opting the FFH Results and the Generality Problem*

Let's suppose that a proponent of a standard theory of justification successfully co-opts these psychological results. So S is justified in adopting the belief that results from her using TB, even when that belief happens to be an I-error (although S does not know it is an error, of course). As we have already noted, the only way that an I-error can be epistemically justified is if we appeal to the reliability of TB: S's belief that p is justified even though it is an I-error because it is produced by TB, which makes more I-errors but fewer overall errors than S's unaided reasoning. But then any theory of justification that sanctions I-errors in this way will be subject to the Generality Problem.

Most philosophers take the Generality Problem to be an objection against reliabilist theories of justification. Such theories make the justificatory status of a belief depend on the reliability of the process that produced it. The Generality Problem arises because for any belief-token, B, and any process-token that produced it, P, there is always more than one way to characterize the process-type of which P is an instance. Some of these process types will be reliable, whereas others will be unreliable. For example, the process token that led to S's I-error (the belief that p) is an instance of at least these two process types:

(a) The reliable process of using TB.
(b) The unreliable process of knowingly and intentionally ignoring most of the relevant evidence (i.e., most of the evidence a regression model or a Bayesian network would have used).

Without some way of deciding which of these process-types to count as *the* one that produced the belief, the reliabilist has no way to decide whether the belief is justified (because it was produced by reliable mechanism (a)) or unjustified (because it was produced by *un*reliable mechanism (b)) (Goldman 1979, Feldman 1985). Here is Richard Feldman's characterization of the problem:

> The fact that every belief results from a process token that is an instance of many types, some reliable and some not, may partly account for the initial attraction of the reliability theory. In thinking about particular beliefs one can first decide intuitively whether the belief is justified and then go on to describe the process responsible for the belief in a way that appears to make the theory have the right result. Similarly, of course, critics of the theory can describe processes in ways that seem to make the theory have false consequences. For example, Laurence BonJour has proposed as counter-examples to the reliability theory cases in which a person believes things as a result of clairvoyance. In his examples, clairvoyance is a reliable process but the person has no reason to think that it is reliable. BonJour claims that the reliability theory has the incorrect consequence that the person's beliefs are justified. He assumes, however, that the relevant process type is clairvoyance. If one instead assumes that the relevant type is "believing something as a result of a process one has no reason to trust" the reliability theory seems to have different implications for these cases (1985, p. 160).

Conee and Feldman describe what must be involved in a solution to the problem, as follows:

> Process reliabilists must solve this Generality Problem. A solution identifies the type whose reliability determines whether a process token yields justification. This type is "the relevant type" for that token. (1998, p. 24)

After reviewing a number of attempts to solve the Generality Problem, Conee and Feldman conclude: "In the absence of a brand new idea about relevant types, the problem looks insoluble. Consequently, process reliability theories of justification and knowledge look hopeless" (1998, p. 24).

In order to co-opt the judgments of the FFH framework, the proponent of a theory of justification must hold that S is justified in accepting I-errors. The only way to get this result is to appeal to the fact that TB is reliable and S knows that it is reliable. *But the Generality Problem rules this out as a consideration that could bestow justification upon this belief token*. As we have already noted, S's belief that p was the result of a reliable process-type (using TB) and an unreliable process-type (intentionally ignoring most of the relevant evidence). The only way S – or any of us – can infer that a belief is justified on the basis of its reliable production is if there is some way to privilege one of these process types as *the one* that determines the belief's justificatory status. But this *is* the Generality Problem. So if the proponent of a theory of justification wants to co-opt the findings of the FFH program, she must face the Generality Problem. The Generality Problem is not just for reliabilists any more.

4.3. *The FFH Framework: Avoiding the Generality Problem*

Any epistemologist who wants to co-opt the prescriptions of the FFH program must face the Generality Problem. For reliabilists, this is not an extra burden: They were already facing the Generality Problem! So perhaps a nice resolution is at hand: The prescriptions of the FFH program give us some reason to be reliabilists. While I have made this suggestion (Bishop 2000), I now think it's probably a mistake. The epistemological framework that grounds the FFH program has a virtue that theories of justification, including reliabilist theories, do not have: It can prescribe on epistemic grounds easier and more reliable reasoning strategies for significant problems (and so can give us the benefits of the research we have considered here) without having to solve the Generality Problem. To see this, recall that this framework evaluates reasoning strategies (not belief tokens) in terms of their reliability and tractability. Suppose that we use this epistemological framework to judge that TB (or the VRAG) is the epistemically best reasoning strategy available to S. So S ought to use TB (or the VRAG); and let's suppose that as a result, S ends up believing that p, and this is an I-error. So far, so good. The Generality Problem arises *only* when we ask: Is S's belief that p justified? But the epistemological framework grounding the FFH program does not address that question. It is not a framework for evaluating individual belief tokens. Unlike traditional theories of justification, the framework that grounds the prescriptions of the FFH program eludes the Generality Problem entirely.

One might be disappointed that this epistemological framework avoids the Generality Problem by failing to address the hard question of whether I-errors are epistemically justified. Note, however, that the framework does tell us that the best reasoning strategy available to S led her to the belief that p. Might there be a way to *extend* this framework so that it evaluates individual belief tokens while avoiding the Generality Problem? Perhaps. But doing that properly would be a very big job – a job for another day.

Notes

[1] I would like to thank Stephen Stich and Shaun Nichols for very helpful comments on earlier drafts of this paper.

[2] Proponents of the FFH program sometimes take FFHs to be decision making heuristics, e.g., the stock-picking heuristic (Borges *et al.* 1999). In fact, Parts II and III of *Simple Heuristics That Make Us Smart* (1999) are called "Ignorance-Based Decision Making" and "One-Reason Decision Making" respectively. This is a mistake, in my view. It is standard practice to evaluate different heuristics against each other in terms of their *accuracy*. This implies that FFHs yield judgments that can be true or false (e.g., which stocks are likely to increase in value). They do not, by themselves, yield judgments about what one ought to do. After all, one might not have the money or the inclination to buy stocks – even if one knows which stocks are most likely to increase in value.

[3] Philosophers might be inclined to call this the cue's reliability. These cross-disciplinary terminological differences can breed confusion. For psychologists, a measurement procedure or instrument is reliable when it gives consistent results across repeated trials. Reliability in the psychologist's sense (consistency) is necessary but not sufficient for reliability in the philosopher's sense (accuracy).

[4] When the training and test sets are not identical, the Profile Memorization Method obviously won't work.

[5] One might reasonably wonder whether at least some of the less-is-more effect is due to the operation of the availability heuristic rather than the recognition heuristic. The availability heuristic involves making frequency judgments in terms of how available information is to memory (Tversky and Kahneman 1974). Goldstein and Gigerenzer distinguish the heuristics by arguing that recognition is all-or-nothing whereas availability to memory admits of degrees (1999, pp. 56–67). It would not be surprising to find that for most Americans, facts about Munich are more easily accessed than facts about Hannover. But in judging Munich to be more populous than Hanover, only 12% of American subjects could have used the recognition heuristic (since all recognized Munich and 88% recognized Hanover). In deciding whether Chicago was more populous than Dallas, only 3% of German subjects could have used the recognition heuristic – and it would have given them the wrong answer (97% recognized Chicago and 100% recognized Dallas). (The recognition rates come from Goldstein and Gigerenzer 1999, p. 55.) It seems plausible that at least some subjects who recognized both cities but who didn't have any well-grounded beliefs about which was more populous made their judgments on the basis of the availability heuristic – on the basis of how easily information about these recognized cities could be dredged from memory.

[6] The notion of "domain specificity" is intuitive, but not well-developed. Is there a principled distinction between rules that are domain specific and rules that are non-domain specific? The rules of logic and probability are prototypical non-domain specific rules. But modus tollens can't be applied to any old problem – even any old logic problem. Ditto Bayes theorem and probability. So why aren't the rules of logic and probability domain-specific? The obvious reply relies on a distinction between form and content: The rules of logic and probability apply to domains that can be defined solely in terms of the formal features of problems. Two worries. First, it is not obvious that a form vs. content distinction is generally clearer than the domain-specific vs. non-domain specific distinction. Indeed, Gigerenzer explicitly rejects the sharp form-content distinction as applied to rules of probability (1996, p. 592). Second, what about Take the Best? It seems to properly apply to a problem with any sort of content, as long as the evidence has features that look quite formal (Martignon and Hoffrage 1999). For proponents of FFHs like Gigerenzer, this would put TB on the wrong side of the divide between domain specific and non-domain specific rules.

[7] Actually, there is some evidence that it was luck. The Borges *et al.* study (1999) was conducted during a historically strong bull (rising) market. Boyd (2001) found that the recognition heuristic, when tested on a bear (falling) market, gave below average returns.

[8] One might object to this project since it (a) presupposes epistemological standards in order to identify the strongest part of these scientific views, and then (b) attempts to extract epistemological lessons from those views. I cannot do justice to this objection here. But it is worth noting that any reasonable epistemological project must presuppose *some* epistemological standards. After all, no one can engage in *high quality* epistemological theorizing while presupposing *no* epistemological standards.

[9] To see this point pressed relentlessly, see Bishop and Trout 2005, 2005a.

[10] One might object to this newfangled Quinean naturalism by insisting that insofar as psychologists are making normative claims, they are no longer doing psychology but philosophy. Therefore, epistemology is still a branch of philosophy. Sidestepping the mind-numbingly pedantic task of defining what philosophy *really* is, I would note that this move does nothing to remove the sting of this version of Quinean naturalism: Epistemology (whatever its home discipline) is something that professional psychologists have been doing very effectively for decades. What's more, on any practical measure, they've been doing it more effectively than professional philosophers.

[11] This project is by no means foreign to philosophers. It is what Descartes was up to in *Rules for the Direction of the Mind*. Others who have pursued it include Mill and Popper.

[12] Why call this an "I-error"? Because intuitively it is an *irresponsible* error – an error one could have avoided if only one had paid "proper" attention to the evidence. Of course, I don't believe that such errors are genuinely irresponsible. Instead, I think that our epistemic intuitions are systematically mistaken about I-errors. (For an argument to this effect, see Bishop 2005.)

[13] I have found some people are inclined to suggest a middle way. Perhaps S is justified in using TB except in those cases in which S comes to believe that TB is making an I-error. This suggestion survives neither reflection nor empirical testing. Recall that we're assuming TB is more reliable than S's unaided reasoning. The "middle way" involves attempting to *improve* the more reliable strategy by allowing the less reliable strategy to "correct" it. That hardly seems reasonable. What's more, the "middle way" has been tested and found wanting. In selective defection studies, subjects are given a simple, reliable rule and allowed to override its judgment. The consistent result of these studies is that subjects are *still* outperformed by the rule, even when they're told that it has been shown to be more reliable than experts (Sawyer 1966, Goldberg 1968, Leli and Filskov 1984). That's not to say that we should never defect from a successful FFH. But we should defect only when it is completely obvious that the FFH has made an error. (For a fuller discussion, see Bishop and Trout 2005, pp. 45–52.)

Works Cited

Barkow, J., L. Cosmides and J. Tooby, eds. *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*. New York: Oxford University Press, 1992.

Bishop, M. "The Autonomy of Social Epistemology." *Episteme* 2.1 (2005): 65–78.

———. "In Praise of Epistemic Irresponsibility: How Lazy and Ignorant Can You Be?" *Synthese* 122 (2000): 179–208.

Bishop, M. and J. D. Trout, *Epistemology and the Psychology of Human Judgment*. New York: Oxford UP, 2005.

———. "The Pathologies of Standard Analytic Epistemology." *Nous* 39.4 (2005a): 696–714.

Borges, B., D. Goldstein, A. Ortmann and G. Gigerenzer. "Can Ignorance Beat the Stock Market." *Simple Heuristics that Make Us Smart*. Ed. G. Gigerenzer, P. Todd and the ABC Research Group. Oxford: Oxford UP, 1999.

Boyd, M. "On Ignorance, Intuition and Investing: A Bear Market Test of the Recognition Heuristic." *The Journal of Psychology and Financial Markets* 2 (2001): 150–6.

Breiman, L., J. Friedman, R. Olshen and C. Stone. *Classification and Regression Trees*. Monterey, CA: Wadsworth, 1984.

Buller, D. *Adapting Minds: Evolutionary Psychology and the Persistent Quest for Human Nature*. Cambridge, MA: MIT Press/Bradford Books, 2005.

Charniak, E. "Bayesian Networks without Tears." *AI Magazine*. 12.4 (1991): 50–63.

Cohen, L. J. "Can Human Irrationality be Experimentally Demonstrated?" *Behavioral and Brain Sciences* 4 (1981): 317–31.

Conee, E. and R. Feldman. "The Generality Problem for Reliabilism." *Philosophical Studies* 89 (1998): 1–29.

Feldman, R. "Reliability and Justification." *The Monist* 68 (1985): 159–74.

Fodor, J. *The Mind Doesn't Work That Way*. Cambridge, MA: The MIT Press, 2000.

Gigerenzer, G. "Content-Blind Norms, No Norms, or Good Norms? A Reply to Vranas." *Cognition* 81 (2001): 93–103.

———. "Fast and Frugal Heuristics: The Tools of Bounded Rationality." *Blackwell Handbook of Judgment and Decision Making*. Ed. D.Koehler and N. Harvey. Oxford: Blackwell, 2004.

———. "How to Make Cognitive Illusions Disappear: Beyond Heuristics and Biases." *European Review of Social Psychology, vol. 2.* Ed. W. Stroebe and M. Hewstone. Chichester: Wiley, 1991.

———. "On Narrow Norms and Vague Heuristics: A Reply to Kahneman and Tversky." *Psychological Review* 103 (1996): 592–6.

———. "Why the Distinction between Single-event Probabilities and Frequencies is Important for Psychology (and Vice Versa)." *Subjective Probability*. Ed. G. Wright and P. Ayton. New York: Wiley, 1994.

Gigerenzer, G., J. Czerlinski and L. Martignon. "How Good are Fast and Frugal Heuristics?" *Heuristics and Biases: The Psychology of Intuitive Judgment*. Ed. T. Gilovich *et al.* Cambridge: Cambridge UP, 2002.

Gigerenzer, G. and U. Hoffrage. "How to Improve Bayesian Reasoning Without Instruction: Frequency Formats." *Psychological Review* 102.4 (1995): 684–704.

Gigerenzer, G. and P. Todd. "Fast and Frugal Heuristics: The Adaptive Toolbox." *Simple Heuristics that Make Us Smart*. Ed. G. Gigerenzer, P. Todd and the ABC Research Group. Oxford: Oxford UP, 1999.

Gigerenzer, G., P. Todd and the ABC group, eds. *Simple Heuristics That Make Us Smart*. New York: Oxford UP, 1999.

Goldberg, L. R. "Simple Models of Simple Processes? Some Research on Clinical Judgments." *American Psychologist* 23 (1968): 483–96.

Goldman, A. "What is Justified Belief?" *Justification and Knowledge.* Ed. G. Pappas. Dordrecht: Reidel, 1979.

Goldstein, D. and G. Gigerenzer. "The Recognition Heuristic: How Ignorance Makes Us Smart." *Simple Heuristics that Make Us Smart*. Ed. G. Gigerenzer, P. Todd and the ABC Research Group. Oxford: Oxford UP, 1999.

Grove, W. M. and P. E. Meehl. "Comparative Efficiency of Informal (Subjective, Impressionistic) and Formal (Mechanical, Algorithmic) Prediction Procedures: The Clinical-Statistical Controversy." *Psychology, Public Policy, and Law* 2 (1996): 293–323.

Kahneman D. and A. Tversky. "On the Reality of Cognitive Illusions." *Psychological Review* 103 (1996): 582–91.

———. "On the Study of Statistical Intuitions." *Cognition* 11: 123–41. Reprinted in *Judgment Under Uncertainty: Heuristics and Biases.* Ed. D. Kahneman, P. Slovic and A. Tversky. Cambridge: Cambridge UP, 1982. 493–508.

Kahneman, D., P. Slovic and A. Tversky, eds. *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge: Cambridge UP, 1982.

Kitcher, P. *The Advancement of Science*. Oxford: Oxford UP, 1993.

Kornblith, H. "The Laws of Thought." *Philosophy and Phenomenological Research* 52 (1992): 895–911.

Leli, D. A. and S. B. Filskov. "Clinical Detection of Intellectual Deterioration Associated With Brain Damage." *Journal of Clinical Psychology* 40 (1984): 1435–41.

Lopes, L. "The Rhetoric of Irrationality." *Theory & Psychology* 1 (1991): 65–82.

Martignon, L. and U. Hoffrage. "Why Does One-Reason Decision Making Work? A Case Study in Ecological Rationality." *Simple Heuristics that Make Us Smart.* Ed. G. Gigerenzer, P. Todd and the ABC Research Group. Oxford: Oxford UP, 1999.

Martignon, L. and K. Laskey. "Bayesian Benchmarks for Fast and Frugal Heuristics." *Simple Heuristics that Make Us Smart*. Ed. G. Gigerenzer, P. Todd and the ABC Research Group. Oxford: Oxford UP, 1999.

Nisbett, R., ed. *Rules for Reasoning*. Hillsdale, NJ: Erlbaum, 1993.

Nisbett, R. and E. Borgida. "Attribution and the Psychology of Prediction." *Journal of Personality and Social Psychology* 32 (1975): 932–43.

Payne, J. W., J. R. Bettman and E. J. Johnson. *The Adaptive Decision Maker.* New York: Cambridge UP, 1993.

———. "Adaptive Strategy Selection in Decision Making." *Journal of Experimental Psychology: Learning, Memory, & Cognition* 14.3 (1988): 534–52.

Piattelli-Palmarini, M. *Inevitable Illusions: How Mistakes of Reason Rule Our Minds*. New York: John Wiley, 1994.

Pinker, S. *How the Mind Works*. New York: Norton, 1997.

Quine, W. V. O. "Epistemology Naturalized." *Ontological Relativity & Other Essays*. New York: Columbia UP, 1969.

Quinsey, V. L., G. T. Harris, M. E. Rice and C. A. Cormier. *Violent Offenders: Appraising and Managing Risk.* Washington, D.C.: American Psychological Association, 1998.

Rieskamp, J. and U. Hoffrage. "When Do People Use Simple Heuristics, and How Can We Tell?" *Simple Heuristics that Make Us Smart.* Ed. G. Gigerenzer, P. Todd and the ABC Research Group. Oxford: Oxford UP. 1999.

Samuels, R. "Evolutionary Psychology and the Massive Modularity Hypothesis." *British Journal for the Philosophy of Science* 49 (1998): 575–602.

Samuels, R., S. Stich and M. Bishop. "Ending the Rationality Wars: How to Make Disputes About Human Rationality Disappear." *Common Sense, Reasoning, and Rationality*. Ed. Renée Elio. Oxford: Oxford UP, 2002.

Sawyer, J. "Measurement and Prediction, Clinical and Statistical." *Psychological Bulletin* 66 (1966): 178–200.

Simon, H. A. *Models of Bounded Rationality*. Cambridge, MA: MIT Press, 1982.

Sosa, E. and D. Galloway. "Man the Rational Animal?" *Synthese* 122 (1999): 165–78.

Stein, E. *Without Good Reason: The Rationality Debate in Philosophy and Cognitive Science*. Oxford: Oxford UP, 1996.

Stich, S. "Could Man Be an Irrational Animal? Some Notes on the Epistemology of Rationality." *Synthese* 64.1 (1985): 115–35.

———. *The Fragmentation of Reason.* Cambridge: MIT Press, 1990.

Todd, P. M. and G. Gigerenzer. "Simple Heuristics That Make Us Smart (with commentaries)." *Behavioral and Brain Sciences* 23 (2000): 727–66.

Tversky, A. and D. Kahneman. "Judgment under Uncertainty: Heuristics and Biases." *Science* 185 (1974): 1124–31.

Vranas, P. "Gigerenzer's normative critique of Kahneman and Tversky." *Cognition* 76 (2000): 179–93.

———. "Single-Case Probabilities and Content-Neutral Norms: A Reply to Gigerenzer." *Cognition* 81 (2001): 105–11.

Weinberg, J., S. Nichols and S. Stich. "Normativity and Epistemic Intuitions." *Philosophical Topics* 29.1 and 29.2 (2001): 429–60.

